



# ESnet

ENERGY SCIENCES NETWORK

## ICNWG, Replication/Versioning, and Data Transfer Working Team

Eli Dart – ESnet

Lukasz Lacinski – Globus

Stephan Kindermann – DKRZ

2018 ESGF meeting

San Francisco, CA

December 6, 2017



U.S. DEPARTMENT OF  
**ENERGY**  
Office of Science



# Outline

- Year In Review
- Current Status
- Performance Capabilities of DTN Clusters

# Year In Review

- Plan:
  - Deploy a test federation at multiple sites
  - Set up data replication within the test federation using Synda
  - Use what we learned to deploy replication in production
- What we did:
  - Deploy test federation: ✓
    - CEDA, DKRZ, IPSL, LLNL, NCI
  - Set up data replication within the test federation using Synda: ✓
  - Use what we learned to deploy replication in production: X
    - Still a work in progress

# Current Status

- Scope reduction: publication of replicas
  - Automatic publication of replicated data out of scope
  - Different sites have different needs/policies
  - Not conflating this with getting the data there
- Different sites in various states of readiness for production deployment
  - Several sites have dedicated DTNs
  - Waiting on integration, final finishing touches
  - Will require some focused effort
- Performance is still low
  - Need some focused work here

# Data Transfers Between Large Sites

- Replication involves moving large amounts of data between sites
- Large-scale downloads for analysis are similar
- Both involve transferring large data sets between large storage systems
  - Not somebody's laptop
  - Tier1 ESGF nodes
  - HPC facilities
  - Support server-side analysis of non-local data sets
  - Support large-scale users who need data for non-canned analyses
- Both download methods use the same data service
  - Follow-on to yesterday's conversation: harden and improve performance of download capability

# Making the list of things to do

- I like the idea of building a solid platform of core services
- Short term action planning: address short term actions to support CMIP6 initially
  - Improve single stream bandwidths to CMIP6 data servers from DTNs
  - Configuration issues at sites
  - Data publication to support download via DTNs
    - Data publication using globus endpoint
- Long term action planning: what has to be prepared to support CMIP6+ in the future
  - Exploit Globus Transfer in replication pipeline
  - Expand DTN deployments to match data scale

# NCAR RDA Data Portal


- Let's say I have a nice compute allocation at NERSC – climate science
- Let's say I need some data from NCAR for my project
- <https://rda.ucar.edu/>
- Data sets (there are many more, but these are two):
- <https://rda.ucar.edu/datasets/ds199.1/> (1.5TB)
- <https://rda.ucar.edu/datasets/ds313.0/> (430GB)
- Download to NERSC (could also do ALCF or NCSA or OLCF)

NCAR's Research Data Archive x

Secure | <https://rda.ucar.edu>

UCAR NCAR Closures/Emergencies Locations/Directions Find People

Hello [dart@es.net](#) [dashboard](#) [sign out](#)

**NCAR UCAR**  **Research Data Archive**  
Computational & Information Systems Lab

*weather • data • climate*

**Go to Dataset:**

[Home](#) [Find Data](#) [Ancillary Services](#) [About/Contact](#) [Data Citation](#) [Web Services](#) [For Staff](#)

**First-time visitor to our site?**  
Please take a [video tour](#) of our home page

**Dataset Search:**  
  [Advanced Options](#)

**Look For Data:**

<b>All Datasets</b>	Variable/Parameter	Type of Data
Time Resolution	Platform	Spatial Resolution
Topic/Subtopic	Project/Experiment	Supports Project
Data Format	Instrument	Location
	Recently Added/Updated	

**Recently Added Datasets:** (within the last 6 months)

- ERA5 Reanalysis Monthly Means
- Daily Gridded North American Snowfall
- ERA5 Reanalysis
- NCAR/MOPITT Reanalysis
- GridRad - Three-Dimensional Gridded NEXRAD WSR-88D Radar Data
- CMIP 5 dataset and code for R parallelization
- Dai and Trenberth Global River Flow and Continental Discharge Dataset
- Dai Global Palmer Drought Severity Index (PDSI)

**Get Help:**

- [Frequently Asked Questions](#)
- [Reset your password](#)
- [A-Z Site Index](#)
- [RDA Users Email List](#)
- [RDA Blog](#)
- [RDA video tutorials](#)
- [Email Us](#)

**From Our Blog:**

- [Accessing RDA OPeNDAP endpoints with authentication](#)
- [All RDA data transfer and processing services restored to production](#)
- [RDA Service Outage July 14-18, 2017](#)
- [RDA web services down for maintenance at 1PM MDT on May 3, 2017](#)

[More blog posts ...](#)

**GLADE Users:**  
Much of the RDA is directly accessible from CISE's [GLOBally Accessible Data](#)




NCAR's Research Data Archive X

Secure | <https://rda.ucar.edu/#!fd?nb=y&b=all&v=Full+List>

UCAR NCAR Closures/Emergencies Locations/Directions Find People

Hello [dart@es.net](#) [dashboard](#) [sign out](#)

NCAR UCAR |  **Research Data Archive**  
Computational & Information Systems Lab

*weather • data • climate*

**Go to Dataset:**

[Home](#) [Find Data](#) [Ancillary Services](#) [About/Contact](#) [Data Citation](#) [Web Services](#) [For Staff](#)

**Look For Data:**

- [Create a New List](#)
- OR --
- Continue Narrowing By:**
  - [Variable / Parameter](#)
  - [Type of Data](#)
  - [Time Resolution](#)
  - [Platform](#)
  - [Spatial Resolution](#)
  - [Topic / Subtopic](#)
  - [Project / Experiment](#)
  - [Supports Project](#)
  - [Data Format](#)
  - [Instrument](#)
  - [Location](#)
  - [Progress](#)
  - [Free Text](#)

**Browse the RDA**

Showing datasets with these attributes: [All RDA Datasets](#) : Full List (680)

Select two datasets and [Compare](#) them. [Reset](#) checkboxes

☐ 1. [Daily Northern Hemisphere Sea Level Pressure Grids, continuing from 1899](#) (ds010.0)

grids contained in this dataset make up the longest continuous set of daily gridded pressure data in the DSS archive. These grids have been ...

☐ 2. [Northern Hemisphere Sea-Level Pressure Grids, continuing from 1899](#) (ds010.1)

continuous time series of monthly gridded Northern Hemisphere sea-level pressure degree latitude/longitude grids, computed from the daily grids in ...

☐ 3. [Northern Hemisphere Daily Sea-Level Pressure Grids for 1880 to 1979](#) (ds012.0)

Northern Hemisphere sea-level pressure data on a 10-degree by 5-degree (36x16) period 1880 to 1979.

☐ 4. [Northern Hemisphere Daily \(and Monthly\) Sea-Level Pressure and 500 mb Height Grids for 1946Jan to 1993Dec](#) (ds018.0)

The gridded daily sea-level pressure analyses in this dataset were produced by the operational models of the U.S. Navy Fleet Numerical Oceanography Center (FNOC). The data are arranged in a ...



# GEOSS Global Atmosphere Forcing Data

ds313.0 ☆

For assistance, contact Chi-Fan Shih (303-497-1833).

## Description

## Data Access

Help with this page: [RDA dataset description page video tour](#)

**Abstract:** GEOSS Atmospheric Forcing data, regridded and prepared as meteorological variables to run CESM and WRF simulations.

**Temporal Range:** 2004-01-02 00:00 +0000 to 2017-10-19 21:00 +0000 (Entire dataset)

✦ [Period details by dataset product](#)

**Updates:** Irregularly

**Variables:**

Surface Pressure

Upper Level Winds

✦ [Variables by dataset product](#)

**Vertical Levels:** See the [detailed metadata](#) for level information

**Data Types:** Grid

**Spatial Coverage:** Longitude Range: Westernmost=180W Easternmost=180E

Latitude Range: Southernmost=90S Northernmost=90N

✦ [Detailed coverage information](#)

**Data Contributors:** [UCAR/NCAR/ACD](#) | [UCAR/NCAR/CGD](#)

**How to Cite This Dataset:**

RIS

BibTeX

Tilmes, S.. 2016. *GEOSS Global Atmosphere Forcing Data*. Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory. <http://rda.ucar.edu/datasets/ds313.0/>. Accessed<sup>†</sup> dd mmm YYYY.

<sup>†</sup>Please fill in the "Accessed" date with the day, month, and year (e.g. - 5 Aug 2011) you last accessed the data from the RDA.

Bibliographic citation shown in [Federation of Earth Science Information Partners \(ESIP\)](#) style

[Get a customized data citation](#)

**Total Volume:** 449.28 GB

**Data Formats:** *netCDF*


**More Details:** View [more details](#) for this dataset, including dataset citation, data contributors, and other detailed metadata

**Data Access:** Click the **Data Access** tab here or in the navigation bar near the top of the page

**Metadata Record:** Display in  format

UCAR NCAR Closures/Emergencies Locations/Directions Find People

Hello [dart@es.net](#) [dashboard](#) [sign out](#)

**NCAR UCAR**  **Research Data Archive**  
Computational & Information Systems Lab *weather • data • climate*

**Go to Dataset:**

[Home](#) [Find Data](#) [Ancillary Services](#) [About/Contact](#) [Data Citation](#) [Web Services](#) [For Staff](#)

 **GEOS5 Global Atmosphere Forcing Data**  
ds313.0 ☆




For assistance, contact [Chi-Fan Shih](#) (303-497-1833).

[Description](#) [Data Access](#)

Mouse over the table headings for detailed descriptions

Data File Downloads		NCAR-Only Access	
Web Server Holdings	Globus Transfer Service (GridFTP)	Central File System (GLADE) Holdings	Tape Archive (HPSS) Holdings
Web File Listing	<b>Globus Transfer</b>	GLADE File Listing	HPSS File Listing

The Research Data Archive is managed by the Data Support Section of the Computational and Information Systems Laboratory at the National Center for Atmospheric Research in Boulder, Colorado. NCAR is sponsored by the National Science Foundation.

Follow us:  Atom  Facebook  Twitter

© 2017, UCAR | [Privacy Policy](#) | [Terms of Use](#) | [Contact Us](#)

# Portal creates a Globus transfer job for us

The screenshot shows the Globus Transfer Files interface in a web browser. The browser's address bar displays the URL: [https://www.globus.org/app/transfer?add\\_identity=32ab4348-9cc6-482a-bc52-240f27...](https://www.globus.org/app/transfer?add_identity=32ab4348-9cc6-482a-bc52-240f27...). The Globus logo is visible in the top left, and navigation links for 'Manage Data', 'Publish', 'Groups', 'Support', and 'Account' are in the top right. Below these, there are links for 'Transfer Files', 'Activity', 'Endpoints', 'Bookmarks', and 'Console'. The main section is titled 'Transfer Files' and features a 'RECENT ACTIVITY' section with three icons. Two endpoint panels are shown side-by-side. The left panel is for 'NCAR RDA dataset archive' with a path of '/ds313.0/'. It shows a file list with a folder '1.9x2.5' and a file 'index.html'. The right panel is for 'NERSC DTN' with a path of '/~/petascale-dtn/project/dtn/src/RDA/'. It shows an empty file list.


Endpoint: NCAR RDA dataset archive  
Path: /ds313.0/

Endpoint: NERSC DTN  
Path: /~/petascale-dtn/project/dtn/src/RDA/

File List (Left Panel):

Item	Type	Size
1.9x2.5	Folder	
index.html	File	258 B

# Submit the transfer job, go about our business

 globus

Manage Data

Publish




Groups ▾


Support ▾

Account

Transfer Files | Activity | Endpoints | Bookmarks | Console

## Transfer Files

RECENT ACTIVITY  1  0  0

Transfer request submitted successfully. Task id: d2776d02-bb6f-11e7-9428-22000a8cbd7d 

Endpoint  ☆

Path  Go

select none ↑ up one folder ↻ refresh list permissions ≡

1.9x2.5 Folder

index.html 258 B

Endpoint  ☆

Path  Go

select all ↑ up one folder ↻ refresh list share ≡

# Data Transfer from RDA Portal – Results

## Activity

☰ Task List



### NCAR RDA dataset archive to NERSC DTN

transfer completed 5 hours ago



Overview




Event Log

Task ID 4f923e48-bb48-11e7-9428-22000a8cbd7d

Owner Eli Dart (dart@globusid.org)

Source NCAR RDA dataset archive   
owner: rda@globusid.org

Destination NERSC DTN   
owner: nersc@globusid.org

Condition SUCCEEDED

Requested 2017-10-27 11:54 am

Completed 2017-10-27 11:58 am

Transfer Settings

- verify file integrity after transfer
- transfer is not encrypted
- overwriting all files on destination

Files 5041

Directories 15

Bytes Transferred 449.27 GB

Effective Speed 1.84 GB/s

Pending 0

Succeeded 5057

Cancelled 0

Expired 0

Failed 0

Retrying 0

Skipped 0

[view debug data](#)

# What's Possible: The Petascale DTN Project

- Built on top of the Science DMZ model
- Effort to improve data transfer performance between the DOE ASCR HPC facilities at ANL, LBNL, and ORNL, and also NCSA.
  - Multiple current and future science projects need to transfer data between HPC facilities
  - Performance goal is 15 gigabits per second (equivalent to 1PB/week)
  - Realize performance goal for routine Globus transfers without special tuning
- Reference data set is 4.4TB of cosmology simulation data

# DTN Cluster Performance – HPC Facilities

## Petascale DTN Project

November 2017  
L380 Data Set

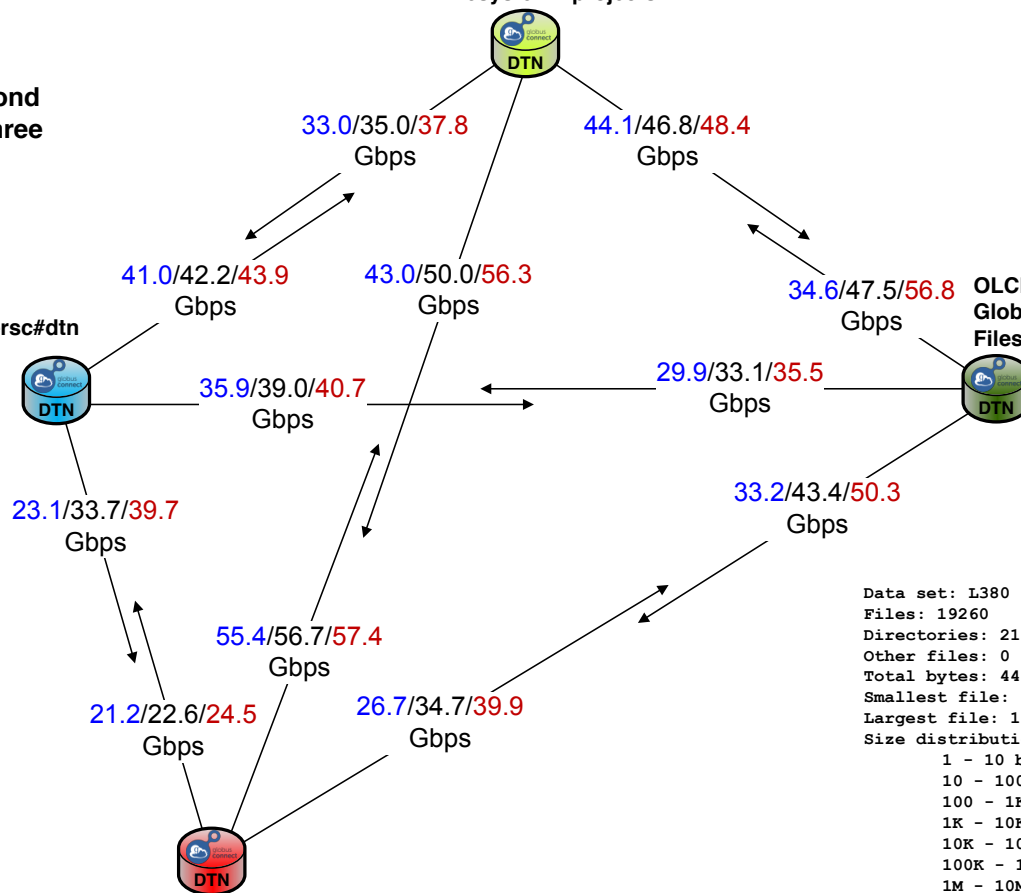
Gigabits per second  
(min/avg/max), three  
transfers

NERSC DTN cluster  
Globus endpoint: nersc#dtm  
Filesystem: /project

ALCF DTN cluster  
Globus endpoint: alcf#dtm\_mira  
Filesystem: /projects

OLCF DTN cluster  
Globus endpoint: olcf#dtm\_atlas  
Filesystem: atlas2

NCSA DTN cluster  
Globus endpoint: ncsa#BlueWaters  
Filesystem: /scratch



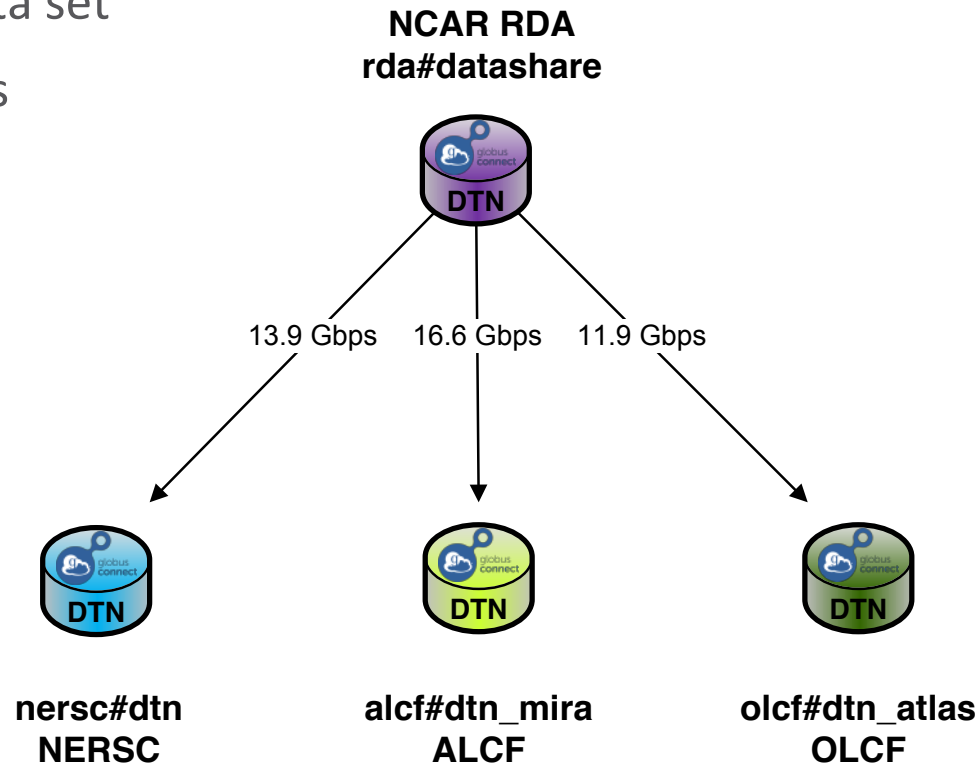
Data set: L380  
Files: 19260  
Directories: 211  
Other files: 0  
Total bytes: 4442781786482 (4.4T bytes)  
Smallest file: 0 bytes (0 bytes)  
Largest file: 11313896248 bytes (11G bytes)  
Size distribution:  
1 - 10 bytes: 7 files  
10 - 100 bytes: 1 files  
100 - 1K bytes: 59 files  
1K - 10K bytes: 3170 files  
10K - 100K bytes: 1560 files  
100K - 1M bytes: 2817 files  
1M - 10M bytes: 3901 files  
10M - 100M bytes: 3800 files  
100M - 1G bytes: 2295 files  
1G - 10G bytes: 1647 files  
10G - 100G bytes: 3 files





# NCAR RDA Performance to DOE HPC Facilities

- 1.5TB data set
- 1121 files



## In conclusion – ESnet's vision:



Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.



# ESnet

ENERGY SCIENCES NETWORK

## Thanks!

Next up: Lukasz Lasinski – Globus

<http://my.es.net/>

<http://www.es.net/>

<http://fasterdata.es.net/>



U.S. DEPARTMENT OF  
**ENERGY**  
Office of Science



# Globus Connect Server v5

- ESGF uses Globus Connect Server v4 which depends on Globus Toolkit
- Support for Globus Toolkit will end in 2018
- Globus Connect Server v5
  - Single port GridFTP (no ephemeral ports)
  - OAuth2 in GridFTP (no more X.509 user certificates or MyProxy)
  - OpenID Connect identity provider
  - HTTPS access to storage
  - Supports multiple connectors in single installation (POSIX, S3, Google Drive)
  - Distributed as Docker containers
  - Released around mid-year 2018
- Globus GUI, Globus CLI, Globus SDK

# ESGF

- Globus Connect Server v5 in ESGF
  - globus-url-copy has to be replaced with Globus SDK/CLI
  - IdPs have to move from OpenID to OpenID Connect (OIDC)
    - planned around mid-year 2018
- Data Transfer Nodes (DTNs)
  - Egress DTN
    - Globus Connect Server v4 with the ESGF authorization callout
    - Documented on <https://earthsystemcog.org>
    - Data sets published with Globus URI (DKRZ, IPSL, JPL, LLNL, NCI, UCAR)
  - Globus Connect Server v5 with the ESGF authorization callout and ESGF OIDC
- Ingress DTN
  - Regular Globus Connect Server v4 or v5
  - Documented on <https://docs.globus.org>



# ESnet

ENERGY SCIENCES NETWORK

## Thanks!

Next up: Stephan Kindermann - DKRZ

<http://my.es.net/>

<http://www.es.net/>

<http://fasterdata.es.net/>



U.S. DEPARTMENT OF  
**ENERGY**  
Office of Science



# ESGF Replication

S. Kindermann, ESGF F2F 2017

# Replication: Discussion Topics

- replication status, problems ..
- Short term action planing to support CMIP6
  - Single stream bandwidth improvments
  - Configuration at sites
  - Data publication and DTNs
- Longer term action planing
  - Globus online exploitation
  - PID exploitation
  - monitoring „dashboard“
  - Full DTN architecture exploitation
  - ..



# Status


Production deployments slowly stabilize, yet unclear overall picture

e.g. last test before departure to the ESGF meeting:

- DKRZ  $\leftrightarrow$  CEDA, IPSL, PCMDI (NCI not reachable)
- Looked completely different some weeks before ..
- $\rightarrow$  besides synda (parallel transfer rates), need to look at single stream transfer rates

gridFTP, 2017-11-30

globus-url-copy: ave(max)



Snet

# Data hubs and data nodes

## Two replication streams to separate

(data hub = larger ESGF replication center)

- data hub  $\leftarrow$  synda  $\leftarrow$  data\_nodes
- data hub  $\leftarrow$  synda  $\rightarrow$  data hub
- (dark data hub  $\leftarrow$  ??  $\leftarrow$  data\_node, data\_hub)




## Need to start coordination discussions:

- Hot spots: identification, synchronization
  - Most accessed variables for CMIP5 available, but ..
- Uncoordinated scenario:
  - Data gets published at a data node
  - All data hubs start replicating
- Data priorities at data hubs ?

### Europe:

- Started IPCC WG1/2 requirements discussion with data hubs at DKRZ and STFC/CEDA
- Evaluation (ESMVal) needs
- Explicit wishes from user groups
- How to coordinate ?
  - Min requirement: make priorities at data hubs transparent – synda sel. files on github .. ?

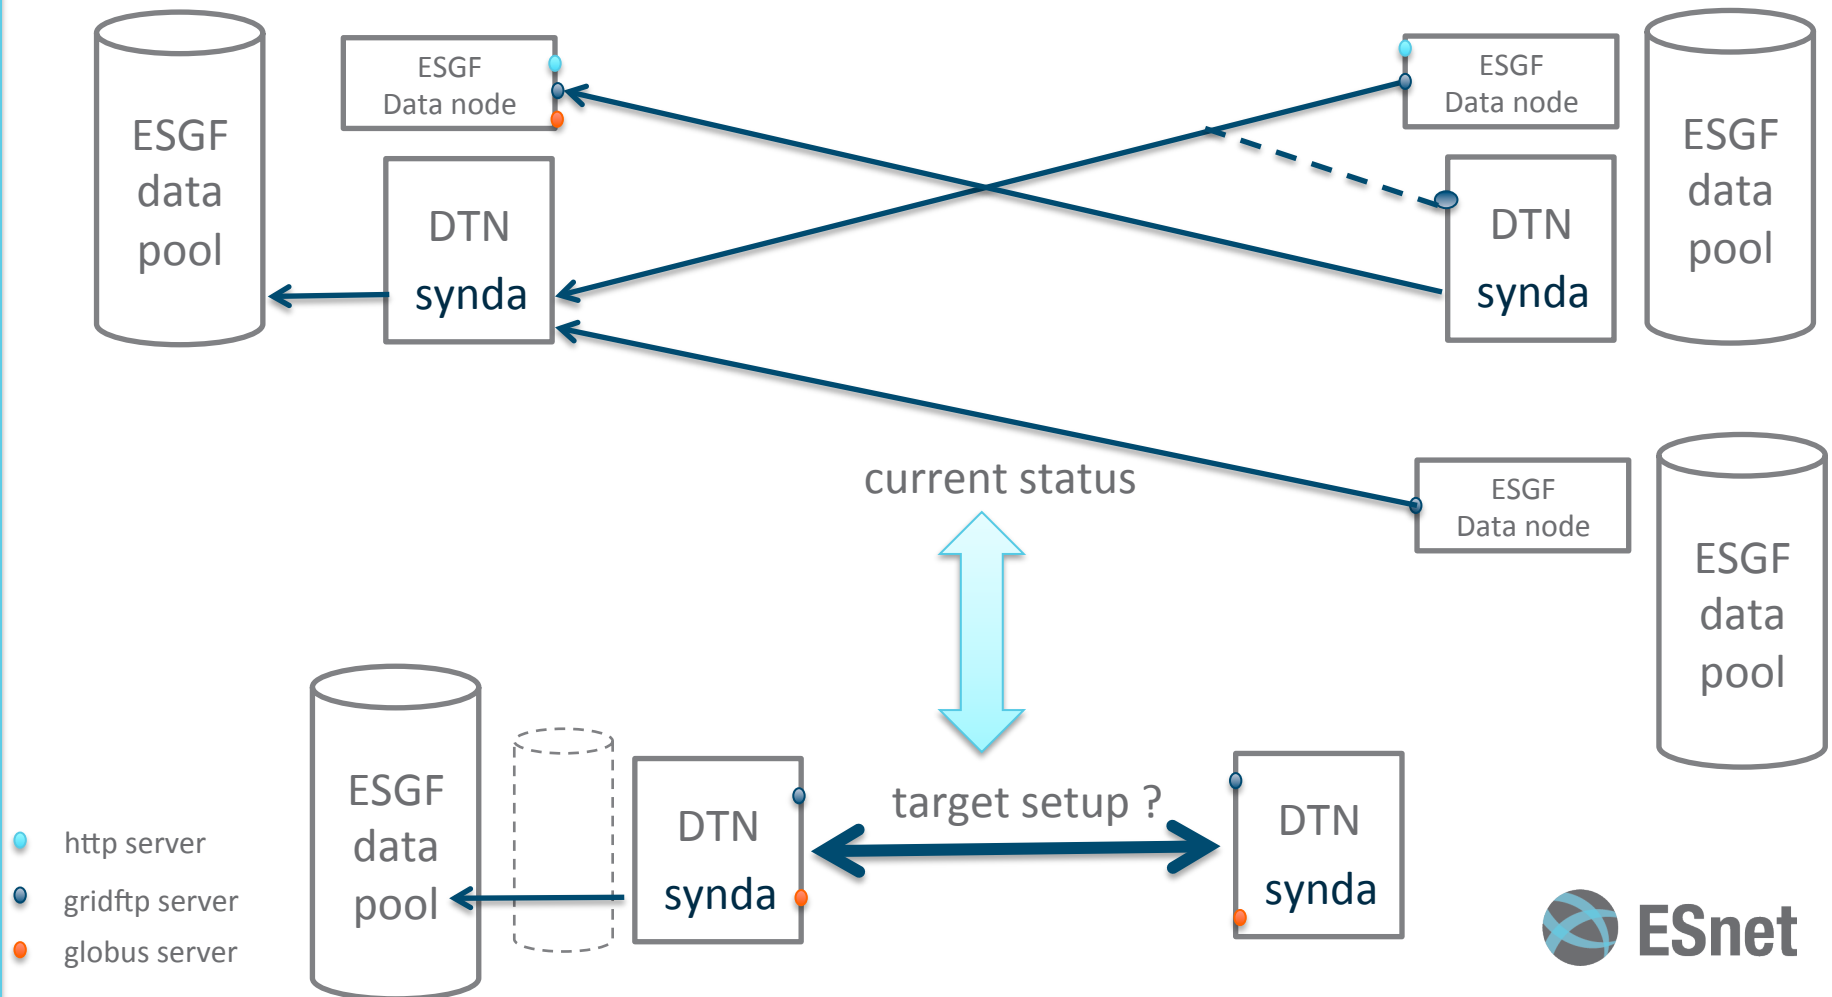
- Ho to coordinate replication ?
  - Publically visible github repos with site specific synda selection files ?, specific datahub repo for coordinated things ?




# Replication: status

Separate gridftp endpoint @ DTN

- Synda config
- endpoint esgf publication






## Problems / Issues

- Installation / deployment
- DTN „visibility“ in ESGF infra
- DTN / firewall / data pool setup


## → Solution approaches

- „esgf“ gridftp/globus installation package
- Agreements with respect to DTN publication
- Best practice collection ..



- Data publication and DTNs
    - Who will provide separate DTN endpoints for data download (besides ESGF datanode gridftp endpoints) ?
- To Do's:**
- Standard instructions in publisher docu
  - Implement consistent DTN endpoint support in infra (naming, index, cog,..)
  - DTN setup instructions (gridftp + esgf AA callout + certificates + (globus) + ..

# Replication: setup alternatives





# ESnet

ENERGY SCIENCES NETWORK

## Thanks!

### Next up: Questions/Discussion

<http://my.es.net/>

<http://www.es.net/>

<http://fasterdata.es.net/>



U.S. DEPARTMENT OF  
**ENERGY**

Office of Science

